

**METHOD FOR TRAFFIC ENGINEERING AND INGRESS ROUTER ADAPTED
TO PERFORM SUCH A METHOD**

The present invention relates to a method for traffic engineering as is described in the non-characteristic part of claim 1, 5 and to an ingress router as is described in the preamble of claim 9.

Such a method is already known in the art, e.g. from *RFC 3270 of the IETF working group, which can be publicly found on the Internet on the web-page <http://IETF.org/rfc.html> and which is titled "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services"*.

10 Therein the basic principles for queuing packets according to their service class, in this document called a Behavior Aggregate and abbreviated with BA, within ingress routers of a packet network such as the Internet, is described. This is supplemented by what is described in another RFC at the same webpage, being RFC3290, in which a BA- 15 classifier and the different queues are further described, the BA-classifier being the device within the ingress source router which determines in which queue a packet will be temporarily stored. Furthermore RFC 3031 from the same webpage describes the concept of having dedicated tunnels to which traffic can flow.

20 State of the art Internet ingress routers are thus adapted to discriminate incoming packets into those following the "normal" IP- route , and those intended to follow the tunnel, in the Internet case being mostly MPLS-tunnels. Moreover, these state of the art routers can further discriminate between the different service classes of the packets 25 such that, per egress interface of such an ingress router, several queues are present. These queues can for instance consist of Diffserv queues in the Internet, and being one per service class, in which all packets intended to be transported over that interface are classified in accordance to their service class. It is important to recognize that in this 30 case no distinction is made between the "normal" IP-traffic, and the

"tunnel" MPLS-traffic, in other words, both traffic type packets are stored in the same queue if these have the same service category.

This state of the art method and ingress router however has the drawback that the MPLS tunnel can be easily overloaded with traffic,
5 since the estimated bandwidth for this tunnel, which is initially communicated by the network administrator and which is reserved for the specific tunnel, is just a prediction which may prove not be very realistic, and which may give rise to congestion problems. In the present situation no traffic engineering solution is thus available for the traffic
10 intended for the MPLS-tunnels.

An object of the present invention is thus to provide such a method for traffic engineering within a packet network, especially for that part of the traffic intended to a specific tunnel .

According to the invention, this object is achieved by the
15 method which further includes the steps as described in the characterising portion of claim 1, and by an ingress router which is adapted to perform these steps as is further described in the characterising portion of claim 9.

In this way, by providing a dedicated tunnel queue per tunnel
20 within the ingress router, and by further shaping the traffic by a dedicated shaper per tunnel queue or per tunnel, a simple method for traffic engineering the traffic intended for the tunnel is provided.

A further characteristic features of the present invention is mentioned in claim 2 and claim 10.

25 By providing, per tunnel, a set of queues, one per service class, differentiation between the service classes for one tunnel is provided. In this case, total tunnel traffic , is shaped, while a set of queues, one per service class is created for the tunnel. The advantage of this is that the total tunnel traffic is limited at a certain shaper rate

while each different service class is still treated separately according to their service class, such as the diffServ characteristics, in the tunnel.

The distinction between the different service classes can be even more further elaborated by providing a separate shaper for each of 5 these queues , as described in claims 3 and 11, such that a traffic engineering tunnel can be further used to engineer multiple service classes. In this case, having one queue and associated shaper per service class of the tunnel allows monitoring and shaping each service class traffic separately.

10 Another characteristic feature of the present invention is described in claim 4 and claim 12.

Thereby one set of queues, each queue of the set pertaining to a different service class, is provided for the traffic for a plurality of tunnels, all tunnels of this plurality pertaining to the same egress 15 interface of this ingress router.

A monitoring device is provided, specifically to control the load or traffic via this dedicated tunnel or the plurality of tunnels, as is described in claims 5, 6 and 13 . This may be performed by periodically measuring the number of packets and their size sent out from the 20 queues, or the number of octets sent out from the queues. On the basis of this monitoring, a comparison can be made with the predetermined reserved bandwidth for the tunnel. This may for instance be performed by comparing the monitored traffic with a predetermined threshold related to this predetermined reserved bandwidth. If this threshold is 25 exceeded, a notification message to the network administrator is generated. The latter can then, based on such a message, increase the reserved bandwidth for the tunnel or the plurality or tunnels, which may in its turn result in calculating a new path for the tunnel or tunnels with this new bandwidth as described in claim 7. Furthermore, this can also

result in providing new shaping parameters by the network administrator to the dedicated tunnel shapers .

In addition to the aforementioned features, the present method is enabled by the network administrator through the sending of 5 a predetermined message to the ingress router , as is further described in claims 8 and 14.

The above and other objects and features of the invention will become more apparent and the invention itself will be best understood by referring to the following description of an embodiment taken in 10 conjunction with the accompanying drawing wherein the figure schematically shows details of an ingress router according to the present invention.

The present invention is used in the field of packet networks, for instance the Internet, wherein, apart from the conventional IP 15 forwarding or next-hop calculation per router, based on the header of each incoming IP packet, also predetermined label switched paths or tunnels are present. In the drawing an ingress router I is depicted to which IP packets may arrive at a number of ingress blades IB1 to IBk. In the drawing an IP packet is depicted which arrives at ingress blade IB1. 20 The determination of whether an incoming IP packet will be forwarded using the conventional IP routing, or will be transferred via the MPLS tunnel, from this ingress router I to an (not shown) egress router , is decided upon within the ingress router, within the FIB-look-up device, denoted FIB. This FIB-look-up device is further adapted to determine 25 which egress blade, and which egress interface of this egress blade of the ingress router I, will be used for sending the packet to. In the figure a situation is depicted whereby egress blade EB1 is selected, and egress interface ITF-1 thereof. In the embodiment depicted in the figure, the ingress router includes n of such egress blades which are coupled, via 30 an internal switch S, to ingress blades IB1 to IBk. Moreover, another

function of this FIB-device is the determination of the tunnel reference, also called LSP-reference, which will be described as well as its use, within a further paragraph of this document.

In addition to the destination information, each incoming IP
5 packet is attributed a predetermined service class, via for instance the
DSCP which is the abbreviation of Differentiated Services Code Point
marker in the header of each packet. For appropriately and adequately
coping with the different bandwidth reserved for these separate classes,
separate queues per class are foreseen per ingress blade , and which
10 are denoted AF1 to AFn, EF, BE and CT. These are respectively the
abbreviation of Assured Forwarding, Expedited Forwarding, Best Effort
and Control Traffic as standardized by the DiffServ working group at
IETF . Traffic pertaining to these different classes will be scheduled and
shaped differently, according to initialized or updated bandwidth and
15 other constraints such as weight, priority (real-time or non-real-time)
constraints. Therefore the incoming packets will first be temporarily
stored in separate queues, per service class in the ingress blade. The
determination of the appropriate queue wherein each incoming packet
is to be stored, is performed in the device BA-IP Classifier, denoted with
20 BA, which determines, based on information within the incoming packet
such as the DSCP marker of the incoming packet, the appropriate
service class.

The information extracted within the BA and FIB-devices, in
the figure denoted with the arrows in dotted lines, is then used within
25 an ingress selector, denoted SELi, in order to determine an internal
switch port such as for instance switch port 1 in the figure, of the
ingress blade, and to determine a specific queue associated to this
internal switch port, towards which the packet will be sent for temporary
storage on the ingress blade.

In the ingress router I depicted in the figure, the incoming packets per ingress blade will be further forwarded towards an appropriate egress interface on an appropriate egress blade. To each of these egress interfaces, and for each egress blade, a similar set of 5 service class queues is foreseen for temporarily storing the incoming packets. In the figure egress blade EB1 is depicted as having three egress interfaces, respectively denoted ITF-1, ITF-2, and ITF-3, and indicated by means of the gray ellipse. One of these egress interfaces, being ITF-3 is only pertaining to classical IP-traffic and has the 10 traditional set of service-class queues, again denoted AF1 to Af_n, EF, BE and CT. ITF-1 and ITF-2 are, apart from the conventional IP-traffic, also adapted to carry tunnel traffic such as MPLS traffic. Two respective tunnels, LSP1 and LSP2, are therefore originating from respectively ITF-1 and ITF-2. For LSP1, L1 is the outgoing MPLS label, for LSP2, L2 is the 15 outgoing label. These are not shown as such in the figure but are important for the further routing of the packet.

In the drawing, the links carrying the IP-traffic are indicated with IP, whereas the tunnels are indicated by means of their tunnel reference.

20 To this purpose, both interfaces are not only coupled to the classical set of queues as described before, but, as an important feature of the present invention, also to at least one dedicated queue per tunnel. This is clear from the figure, where interface ITF-1 is not only coupled to a set of queues similar to the one of for instance interface ITF-3, but is also coupled to a dedicated tunnel queue denoted QLSP1. Similarly, to ITF-2 is coupled a dedicated tunnel queue QSLP2, 25 apart from the conventional set of queues.

The determination of the queue on the egress blade, where each packet will be temporarily stored, is performed in several steps.

30 Firstly a LSP-ref-check device, denoted LSP-r on the figure, is adapted

to check whether the incoming packet on the egress blade EB-1 has to follow the classical IP forwarding, or an MPLS-tunnel. If classical IP, an egress selector, denoted SELe, having similar functionality as the ingress selector SELi on the ingress blade, determines the appropriate egress 5 interface and queue thereof, on the basis of the service class and egress-interface reference . This information is for instance derived within the SELe device itself from a special packet header which was added in front of the IP-packet, within an encapsulating device (not shown on the figure) ingress blade . This special header contains 10 internal parameters like service class, egress-blade reference, egress-interface reference, LSP reference etc. , which were earlier determined within the ingress blade of the ingress router, within devices such as BA and FIB.

In case the LSP-r device finds out that the packet is to be sent 15 via an MPLS-tunnel, an out-segment table, denoted OST in the figure, is used to determine the appropriate storage queue on the basis of the tunnel reference or tunnel label , extractable from the special packet header.

In another embodiment (not shown on the figure) even a set of 20 queues, one for each service class, for one or more tunnels pertaining to the same egress interface, is present. This allows to further differentiate the MPLS traffic across the different service classes within the same tunnel, in case a set of queues exists for one tunnel, or within a group of tunnels , in case a set queues exists for a group of tunnels pertaining 25 to the same egress interface.

In the shown embodiment, whereby for each MPLS tunnel, one queue was foreseen , also one associated shaper is present in the egress blade. These are respectively denoted SLSP1 and SLSP2, and will then adapt the traffic for the respective MPLS tunnel LSP1 and LSP2, in

accordance with the reserved bandwidth such as the Peak Information Rate configured for the queue.

It can be further remarked that such shapers may also be present, although not shown in the figure, for each IP queue.

5 For the tunnel or MPLS shapers, embodiments whereby the Peak Information Rate of the separate shapers is set to the reserved bandwidth of the tunnel, are possible. However other shaper devices may be provided, where other traffic parameters ,determined initially by the network administrator, are used . Since these shapers are well 10 known to a person skilled in the art, such shapers will not be further discussed into detail.

To determine in which queue an MPLS packet will be stored, the OST is extended with the queue-reference. In another embodiment, in case of several tunnel queues per tunnel, according to their service 15 class, an OST-table with also only one queue reference added can be envisaged, whereby this extra reference will then be a reference to a tunnel-queue-block. At the entry of the queue-block it can then be further determined which actual queue will be taken for the storage of the packet. However, the out-segment table could as well be updated 20 with the actual queue reference.

The OST-table, as depicted in the figure, uses the tunnel reference as index to this table, and includes entries such as the outgoing label of the tunnel (L1 or L2), the egress interface (ITF-1 or ITF-2) and the queue reference (QLSP1 or QLSP2).

25 In the figure each egress blade further includes a monitoring device. However other embodiments may include monitoring devices per queue, or per egress interface. The function of such a monitoring device is to monitor the traffic via the tunnels. This may be performed by monitoring the queues attached to any of the egress interfaces of an 30 egress blade. To this purpose, this device is adapted to monitor the

amount of the traffic sent from the queue, for instance by checking the occupation of each queue, and to compare this with a predetermined threshold related to an initial reserved bandwidth for this tunnel . Furthermore such a monitoring device is further adapted to generate a 5 message to the network administrator in case of overflow conditions. Thus the occupancy of the respective tunnel queues will be monitored, and in case of overflow, a message will be generated to the network administrator, indicative of traffic problems such as congestion. The network administrator (not shown on the drawing) can then adapt the 10 tunnel, this generally implying determining a new "tunnel path", possibly including the selection of a new egress blade, and a new egress interface. This means that this information again has to be provided to the FIB classifier. Also the shaping device attributed to the tunnel has to be informed since traffic will now have to be differently shaped.

15 An additional feature of the method of the present invention is that the network administrator can enable this method, thus can enable the feature of having the separate queue per tunnel. To this purpose a message is sent (not shown in the drawing) from the network administrator, for instance by means of the Simple Network 20 Management Protocol, abbreviated by SNMP-protocol, wherein a new tunnel configuration object indicates or orders the ingress router to enable such a dedicated queue for a tunnel. This is usually performed by means of an additional management object of a so-called MIB, being the abbreviation of Management Information Base . However, other 25 means of communication are possible , for instance by using the CLI Command Line Interface. The ingress router is then also adapted to receive such a message, and to extract from its contents the indication whether or not to enable such a separate tunnel queue per tunnel, and to enable the queue in the requested case. In another embodiment, 30 where several queues per tunnel, pertaining to different service classes,

are possible, this message from the network administrator to the ingress router may as well contain details about the enabling of the plurality of queues per tunnel. Similarly, in these embodiments the ingress router is then further adapted to extract from the contents of 5 this message whether to enable the queues or not, and accordingly perform so.

While the principles of the invention have been described above in connection with specific apparatus, it is to be clearly understood that this description is made only by way of example and 10 not as a limitation on the scope of the invention, as defined in the appended claims.